

Health Reporting Mechanism for Inter-Network Gateway

TECHNICAL FIELD

[0001] The present invention relates generally to communication networks, and more particularly to a health reporting mechanism for inter-network gateways.

BACKGROUND

[0002] A telecommunications network supports voice and data communications between customers. What is typically viewed as a single network, however, can actually be a series of separate networks, many times owned and operated by different companies. Oftentimes, a "single" data communications network will have portions that utilize different technologies. For example, one part of the network may be based upon frame relay technology while another part of the network is based upon asynchronous transfer mode (ATM) technology. A gateway is a network component that bridges these different portions.

[0003] One goal of a telecommunications service entity is to maintain the level of quality of the network. For example, entities presently provide for the presentation and dissemination of customer account and network data management information to their customers by, for example, enabling customers (clients) to connect to the entity's application servers to access their account information. The requests are processed by the entity's application servers, which retrieve the requested customer information from one or more databases, process and format the information for downloading back to the client.

[0004] As an example, larger telecommunications inter-exchange carrier enterprises provide management and performance information relating to circuits comprising a customer's

broadband network, including web servers as an example. Such network management information generally includes details of network use and performance such as, for instance, real time status and alarm information, near real time performance data, usage statistics, SNMP data, etc. For example, the carrier could provide a system that monitors all aspects of web server health from CPU usage, to memory utilization, to available swap space so that Internet/Intranet networks can increase their hit rate and reduce Web server management costs. Software processes can generate alerts based on process health, connectivity, and availability of resources (e.g., disk usage, CPU utilization, database availability).

[0005] Such health reporting mechanisms have been implemented within a particular network. For example, a system has been implemented to generate health reports for a Frame Relay network. The system polled Frame Relay switches to obtain information on parameters such as CPU and memory utilization. This information was analyzed and provided to support staff on a regular basis so that it could be studied for conditions that might lead to a network failure or other instability. Similar health reporting has also been implemented in a Private IP (PIP) network.

SUMMARY OF THE INVENTION

[0006] In accordance with a preferred embodiment of the present invention, a method for using a data communications network includes receiving at a gateway device a first communication from a first network that is addressed for a network element of a second network. In this case, the second network is based on a different technology than the first network and the gateway device is a layer 3 gateway. The first communication is then transmitted from the gateway device to the second network. The gateway device also receives a second communication from the second network that is addressed for a network element of the first network. This second communication is transmitted from the gateway device to the first network. The gateway device is also periodically polled to obtain operating parameters related communications between the first and second networks. These operating parameters are analyzed and a health report is generated based upon analysis of the operating parameters. The health report is related to the gateway device(s), and possibly the first network and the second network.

[0007] In another aspect, the present invention provides a method of monitoring the stability of a network. An inter-network gateway is periodically polled to collect data related to the inter-network gateway. For example, the data related to at least one of a flowcache, a virtual private routed network, or an internet key exchange security association. The data is processed to generate a number of parameters and a report is generated based on these parameters. The report can then be automatically transmitted, e.g., transmitted without human intervention.

[0008] In yet another aspect, the present invention provides a computer program for use in a system for monitoring the stability of a data communications network. The computer program is

operable to periodically gather information related to the network and provide a report related to the gathered information. In the preferred embodiment, the computer program includes computer program code for automatically, periodically polling a plurality of inter-network gateways to collect data related to the inter-network gateway, code for processing the data to generate a number of parameters, code for generating a report based on the parameters, and code for automatically transmitting the report.

[0009] An advantage of a preferred embodiment of the present invention is that it provides a tool to predict network instabilities by proactively monitoring selected parameters. Network performance reporting tools available on the market today, such as software developed by Quallaby Corporation, CrossKeys Systems Corp. and Concord Communications, Inc., provide information on network availability, latency, and throughput to support customer service level agreements (SLAs). Performance-based tools, however, have been proven to be ineffective as a predictor of events that could result in a network brownout or outage.

[0010] Another class of software, network fault management tools, such as NetCool™ and NetExpert™ report alarms. Alarms are generated when a network element, such as a switch, exceeds a pre-defined threshold. Alarms indicate that a problem exists in the network that requires intervention. As such, it is indicative of a problem that is already occurring.

[0011] The health reporting mechanism of the preferred embodiment, on the other hand, detects events and conditions that could lead to future network problems. Frequently, these events and conditions do not generate alarms and would otherwise go undetected until they become service impacting. Comparing these tools to human health, a performance tool would measure the distance and speed of a runner, a fault management tool would detect if he has pulled a muscle, and a health reporting tool would monitor blood pressure, pulse, etc.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

[0013] Figure 1 is a depiction of an inter-network gateway, the networks connected through it, and a health reporting system;

[0014] Figure 2 is a block diagram of one example of the health reporting system of Figure 1;

[0015] Figure 3 is a flow diagram of the report generation process; and

[0016] Figure 4 is a flow diagram of the analysis of the report, and disposition of the findings.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[0017] The making and using of the presently preferred embodiments are discussed in detail below. It should be appreciated, however, that the present invention provides many applicable inventive concepts that can be embodied in a wide variety of specific contexts. The specific embodiments discussed are merely illustrative of specific ways to make and use the invention, and do not limit the scope of the invention.

[0018] The present invention will be described with respect to preferred embodiments in a specific context, namely a network of gateways that couple data communications networks of different technologies. The invention may also be applied, however, to other networks such as voice networks.

[0019] In one aspect, the present invention provides a health reporting mechanism for an inter-network gateway, e.g., a layer 3 gateway. For example, in June 2003 MCI introduced a new network-based service designed to securely and seamlessly connect remote and traveling workers to their existing corporate networks via the Internet. This network, referred to as the Secure Interworking Gateway (SIG), provides enterprises with access to their existing frame relay, ATM, Private IP and IP VPN networks via the Internet. The gateway automatically authenticates the connection, requesting user name and password confirmation from the edge server. Once a user is granted access, communications are routed to a customers' data network via a Permanent Virtual Circuit or an IPSec tunnel.

[0020] The preferred embodiment of the present invention provides a health reporting mechanism that can be utilized with a gateway network such as the SIG. In one embodiment, a combination of software scripts (code), reports, processes, and benchmarks allow network

engineers to accurately and efficiently assess the health (stability) and performance of any size network. This tool identifies and reports undesirable network conditions that could result in a “brownout” (i.e., significant degradation of performance) or total outage. In one aspect, it is a predictive tool that is designed to spot network problems before they can become service impacting.

[0021] Aspects of the present invention are particularly useful with an inter-network based gateway (or a network of such gateways). Accordingly, such a network will be described first including exemplary services that can be provided by the network. A preferred mechanism for monitoring the stability of such a network will then be described.

[0022] Referring first to Figure 1, a data communications network 100 includes a number of network portions 102-106 that can be interconnected by an inter-network gateway 110. Each of the network portions 102-106 utilizes a different technology. For example, network 102 is a frame relay (FR) network, network 104 is an asynchronous transfer mode (ATM) network, network 106 is a private internet protocol (PIP) network, and network 108 is a (internet protocol virtual private network (IPVPN)).

[0023] While illustrated with a single cloud, each network 102-108 can comprise multiple networks, e.g., from a hardware or a billing perspective. In addition, these four network types are provided as examples. More or less than four types, including these or other networks, can utilize aspects of the present invention. Each network includes a number of network elements, server 118 being provided as an example.

[0024] A discussion of a network of the type shown in Figure 1 is provided in the White Paper by Robert Eppich entitled "Bridging the Great Divide Between Public and Private

Networks," Nov. 2003, which paper is available at <http://global.mci.com/us/enterprise/insight/whitepapers/pdf/SIG.pdf> and which is incorporated herein by reference.

[0025] Network portion 102 is a frame relay network. Frame relay technology is based on the concept of virtual circuits (VCs). VCs are two-way, software-defined data paths between two ports that act as private line replacements in the network. Permanent virtual circuits, or PVCs, are set up by a network operator via a network management system. PVCs are initially defined as a connection between two sites or endpoints. New PVCs may be added when there is demand for new sites, additional bandwidth, alternate routing, or when new applications require existing ports to talk to one another.

[0026] Frame relay is a synchronous protocol where data is carried across a communications line in frames that are similar in structure. In a frame relay frame, user data packets are not changed in any way. Frame relay simply adds a two-byte header to the frame. The frame relay header contains a 10-bit number called the Data Link Connection Identifier (DLCI). The DLCI is the frame relay VC number (with local significance), which corresponds to a particular destination. The frame relay switches utilize routing algorithms such as open shortest path first (OSPF) to determine the optimal path for the PVCs. The DLCI addresses are used by the intermediate frame relay switches to uniquely identify the PVCs and determine the optimal path.

[0027] ATM technology can be used in the portion 104 of communications network 100. Asynchronous Transfer Mode technology (ATM) is also a layer 2 networking technology based on the concept of using VCs that are set up by a network operator. Unlike frame relay's variable-length packets, the ATM protocol uses fixed-length packets (cells) to transport user data across the network. It is the use of these fixed-length cells that allows ATM to support a range of applications and traffic types. Cells are statistically multiplexed and network capacity is

dynamically allocated based on the real-time needs of the applications supported. These cells consist of 48 bytes of data payload and a five-byte header containing the addressing information required for information delivery. This header contains Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) addressing information performing a similar function to the frame relay DLCI.

[0028] Private IP service, as shown by cloud 106, is based on multi-protocol label switching (MPLS) technology. MPLS enables networks to take advantage of the best of IP, ATM, and frame relay by allowing the integration of layer 2 switching (ATM and frame relay, for example) and layer 3 routing (IP). The MPLS signaling protocols support and create labels required to move the traffic across the network. The labels identify the end-address destinations of the network traffic as well as quality of service (QoS) information for the prioritization of traffic across the network. The QoS information is initially applied by the CPE router to the Type of Service (TOS) byte of the IP packets and this information is copied into the MPLS labels by the provider MPLS switches. This QoS capability allows for the transmission of both real-time applications such as voice and video along with data traffic across the same network infrastructure.

[0029] MPLS-enabled Private IP networking offers full IP routing capabilities at the edge of the network. The customer premises equipment (CPE) peers with the MPLS-enabled switch, and IP routes are exchanged using static routes, RIPv2, EBGP, or OSPF. A virtual router is defined in the software of the MPLS-enabled switch, which is unique to the customer, and the customer CPE peers with this virtual router using a layer 2 encapsulation protocol. IP data is encapsulated by the CPE in layer 2 ATM or frame relay for transport to the MPLS-enabled core. At the MPLS switch, the layer 2 encapsulation is stripped and the layer 3 IP packet is used to make

routing decisions. Based upon the destination IP address, this MPLS-enabled switch knows where the packet will leave the switched network and uses that as its target destination. Each IP packet is then encapsulated in an MPLS header and all further switching within the backbone network is performed based on label swapping. In the final MPLS-enabled switch, an IP lookup determines the outbound port destination of the packet, the MPLS labels are removed, and the IP packet is once again encapsulated in ATM or frame relay for delivery to the destination node.

[0030] Private IP service utilizes frame relay and ATM PVCs for access into the MPLS-based core network. This allows for a simple migration from an existing frame relay or ATM network. These PVCs peer with the edge MPLS switch and not the destination router; frame relay and ATM is simply used as transport mechanism to the MPLS edge router. Additional benefits include inherent any-to-any connectivity, class of service offerings, and seamless interworking between frame relay and ATM similar to FRASI described earlier.

[0031] While layer 2 communication is handled from location to location via a PVC, Internet protocol (IP) requires each device connected to the Internet to be identified by a unique number, the IP address. Since there is no permanent circuit required for locations and devices to communicate with one another, this allows for a connectionless network. It is this connectionless-oriented network that has allowed for the rapid availability of the Web and its associated solutions. The Internet 112 is the largest public network, as it is a culmination of many service providers' facilities connected in a hierarchical fashion. Anyone can freely participate in this network as long as they register themselves with a unique IP address. One concern with using this infrastructure is the lack of security; again it is open to anyone in the world. Security measures have been developed to allow for secure transmission across this architecture which involve the encryption of the data prior to transmission across secure

“tunnels” and the subsequent decryption upon data receipt. IP VPN 108 can be thought of as a “network” based on these tunnels.

[0032] Tunneling refers to the creation of a secure temporary path over an inherently unsecured network such as the aforementioned public Internet. While there are a number of tunneling protocols (L2TP, PPTP, L2F), the most prevalent in VPN deployments is the layer 3 tunneling protocol suite, IPSec. The IPSec protocol suite enables authentication, confidentiality, and integrity between systems. Here it is important to point out that IPSec does not authenticate users, but authenticates devices. The IPSec tunneling process is established in three phases; determine whether IP communications require IPSec, negotiate and then establish the secure communications, and transmit the data.

[0033] The IPSec protocol suite can provide for data origin authentication, anti-replay, integrity, and confidentiality. Most implementations of IPSec accomplish this using Internet Key Exchange (IKE) and Encapsulating Security Protocol (ESP). A prerequisite to an IP packet being secured by IPSec is that a Security Association (SA) must exist. This SA may be created manually or dynamically. IKE is used to create them dynamically on behalf of IPSec and requires that the IPSec peers first authenticate themselves to each other and then establish a shared key for encrypting and decrypting data. Once an SA is established, ESP is used to perform the data authentication, antireplay, integrity, and confidentiality. It does so by using a combination of cipher and authentication algorithms and then inserting a protocol header into the IP datagram that provides the information required to perform these functions.

[0034] The other network illustrated in Figure 1 is the public Internet 112. In the illustrated embodiment, the public Internet 112 denotes the worldwide collection of interconnected networks that uses Internet Protocol to link the large number of physical networks into a single

logical network. Physically, the Internet is a huge, global network spanning countries around the world and comprising a great number of academic, commercial, government, and military networks.

[0035] Inter-network gateway 110 is provided to allow communications (interworking) between each of the network portions 102-108. Inter-working allows users to share information privately and securely across a variety of physical network topologies. Gateways 110 are responsible for logical connection termination, authentication, security, and protocol conversion. One example of a gateway 110 is the Shasta broadband service node (BSN) available from Nortel Networks.

[0036] In the preferred embodiment, the gateway 110 is a layer 3 gateway. In other words, the gateway routes data from one network to another based on the layer 3 address, regardless of the arbitrary layer 2 encapsulations at the edges of the individual networks. This device can be distinguished from a network-to-network interface (NNI) that splices connections between edges of different networks of the same layer 2 technology (e.g., frame relay or ATM) and from trunks that transfer data within the core of a single network during long haul communications.

[0037] In the preferred embodiment, the gateway 110 has user-to-network (UNI) links to frame relay and ATM networks 102 and 104 for PVCs to the customer's private frame relay, ATM, or Private IP WAN network. The gateway PVC endpoint terminates on a static virtual interface configured on the customer's VPRN (virtual private routed network). The other PVC endpoint terminates on a router on the customer's enterprise network, or in the case of PIP, on the customer's virtual router. These point-to-point links appear as static subscribers to the VPRN. As an example, they can be configured as IP.252 subnets (/30) and can run RIPv2 or carry static routes for reachability to the customer's WAN and beyond. IP VPN Dedicated

networks are connected to their configured VPRNs by virtual interfaces terminating IPSec tunnels through the IP trunks.

[0038] The VPRN functions as a conventional IP router, with awareness only of the data arriving over the active virtual interfaces, so that there is no constraint on the customer's IP addressing plan. It builds a route table with static and/or dynamically generated routes, and forwards datagrams to the virtual interface indicated by the route table. In this way, traffic can be routed between many remote clients and destinations on the enterprise network.

[0039] The inter-network gateway 110 allows the service provider to provide a number of services between network portions 102-108 and the public Internet 112. For example, remote access occurs when a remote user 116 connects to a server 118 via the Internet 112. An example of remote access occur when a corporate employee wants to connect to the company's internal network while traveling or working from home.

[0040] To activate remote access, the user 116 launches a connection to the Internet 112, e.g., through an Internet service provider 114, and communications from the user would be routed to the gateway 110, which has a trunk to the Internet. The gateway 110 will authenticate the user 116 and create a bridge between the access device 116 and the corporate network 118. In this respect, the gateway 110 is generic on the access side, since it can connect to the Internet, and can have proprietary connections on the private side. These connections allow communication with any portion of the network 100 regardless of the underlying technology.

[0041] In a practical application, the gateway 110 is implemented as a network of gateways. For example, the gateway 110 may include many (e.g., 10-100) physical boxes that can be located at numerous places throughout the network, within the country and around the world. In some instances, the gateway 110 will be located in the same physical location or terminal as the

frame relay 102 or ATM 104 equipment. Due at least to the Internet connection, the device 110 is logically reachable from anywhere in the world.

[0042] The gateway 110 enables a network-based remote access solution by serving as a concentrator switch that is a shared device hosted in the provider's network. As noted above, this shared device 110 includes multiple virtual routers and each customer 118 is provisioned to a unique virtual router through their private network connection. Access to this shared switch 110 is limited to the provider and security policies are implemented to ensure that privacy is maintained between all customers. This feature is attractive to customers with existing private network infrastructures that do not want to procure secondary connectivity directly into the public Internet. Since the provider manages the network-based gateway 110, the customer can extend a PVC or dedicated IPSec tunnel from their host site across their existing private infrastructure into this network gateway 110. This private connection terminates onto a virtual router in the network-based gateway assigned exclusively to that customer.

[0043] The customer's virtual router also has connectivity into the public Internet 112. With this service, the customer purchases a PVC into the gateway 110 as opposed to purchasing an additional Internet connection and concentrator equipment as would occur in a CPE-based remote access solution. This is the optimal low-cost solution for customers with existing private network connectivity and sufficient available bandwidth at their host port to support the additional dial user traffic by extending a PVC or IPSec tunnel to a virtual router in the network-based virtual gateway.

[0044] Another feature that is provided by the gateway 110 is a firewall. The firewall features within the gateway 110 will intercept packets to verify conformance to a rule set before forwarding is allowed. This task can be accomplished without additional physical connections.

The firewall solution can be advantageous since the customer 118 would not be required to purchase a firewall at the customer premise. Rather, the provider builds the PVC into a network-based firewall, which is hosted and managed within the provider's resilient network 100. The network provider is then responsible for the management and maintenance of the hardware and software. Preferably, the customer will have the ability to access the firewall via a secure GUI in order to manage and set policies for their particular security requirements.

[0045] When operating a network 100 it is important to avoid network crashes. One method to accomplish this goal is to have a system in place to monitor and report on significant network parameters and statistics. One goal of this system is to anticipate network instabilities so that they can be corrected before the network performance is adversely affected.

[0046] Accordingly, a health reporting system can be used to gather, analyze and distribute information about events and conditions in the network 100 that may affect network stability and performance. The health report can be transmitted (e.g., e-mailed) daily to network engineering and operational support staff with responsibility for ensuring the health and performance of the network. With this information, the staff can then make informed judgments and take proactive steps to avoid or reverse complications, in the same way that a doctor would respond to a report that a patient has high blood pressure or cholesterol.

[0047] In one embodiment, the report contains information about network events such as card toggles, and about network statistics such as CPU and memory utilization for the numerous processors, and counts of significant logical objects such as connections, VPNs, and IPSec tunnels. It provides comparative benchmarks for a similar "healthy" network.

[0048] The health reporting can be accomplished by a computer system 120 that is connected to the gateway 110 through a control network (not explicitly shown) or through any of

the networks 102-108 or 112. The particular configuration shown in Figure 1 is merely exemplary of one particular configuration. In one example, the computer system 120 can be a PC running software that polls the gateway to gather specified information and then analyzes this information. This computer can be implemented as a single box or can be implemented over a number of distributed machines, each machine performing a portion of the tasks.

[0049] Figure 2 illustrates an example of a computer system 120 that can be used to implement the health reporting mechanism described herein. In this example, the system 120 includes a processor 120, which runs the software that implements the health reporting mechanism. Further detail regarding the software is provided in the flowchart of Figure 3. In the preferred embodiment, the processor is a microprocessor, e.g., a Pentium™ class chip available from Intel, an Athlon™ class chip available from Advanced Micro Devices, a Sparc™ class chip available from Sun Microsystems, or the like. In other embodiments, other controllers or processing units can be used.

[0050] A number of components are shown attached to the processor 122. For example, random access memory (RAM) 124 can be used to store the program code and working data for the processor. This memory is preferably implemented with a dynamic random access memory (DRAM) in combination with a faster cache memory, e.g., implemented through static random access memory (SRAM) technology. Other memories, such as electrically erasable programmable read only memories, can alternatively be used.

[0051] Long term memory 126 is used to store program code, and possibly other information, while the computer is not being used. As such, memory 126 is preferably a non-volatile memory. In the preferred embodiment, memory 126 is a hard disk drive. In other embodiment, memory 126 can be implemented with an optical drive (e.g., CD-ROM, or DVD),

flash memory (or other non-volatile semiconductor memory), or floppy disk drive.

Combinations of the various types of memory could also be used. The memory 126 could be eliminated if program code and operating system are accessed through the NIC 128.

[0052] Network interface card (NIC) 128 is used to provide a connection to elements outside the system 120. In Figure 1, the system 120 is shown with a connection to the gateway 110. This connection could be made through a NIC 128. In the case where the system 120 is distributed over a number of separate boxes, each box might also have a NIC 128 to facilitate communications.

[0053] Finally, input/output (I/O) block 130 is provided to show that information must be entered into and received out of the system 120. For example, the input portion could include a keyboard and a mouse while the output portion includes a display and a printer. As with the hard drive, the I/O block 130 could be eliminated if all user communications with the system 120 are performed through the NIC 128.

[0054] Figure 3 provides a flow chart 200 showing the operation of a preferred embodiment of the present invention. The analysis system (not shown) communicates with inter-network gateways 110 via an Internal Data Network (not shown) that is used for network management and which is not accessible to customers. As discussed above, the inter-network gateway 110 typically comprises a number of gateway devices, each of which will be polled by the system. Information gathered from the gateways 110 will then be stored, analyzed and summarized in a health report that can be provided to maintenance staff. The functions of the system will be described with respect to Figure 3 and the functions of the staff will be described with respect to Figure 4.

[0055] Referring now to Figure 3, a poller is initiated at periodic intervals as shown by block 210. This initiation is preferably automated, i.e., it begins periodically without human intervention. In the preferred embodiment, the system comprises computer software operating on a UNIX™ operating system and the poller is a software application. In particular, the poller is initiated using a CRON utility. It is understood, however, that any periodic processor can be implemented to automatically run the process at specified intervals.

[0056] In the preferred embodiment, not all operating parameters are gathered at the same intervals. For example, some parameters are gathered every fifteen minutes, others every hour and yet others once a day. These time intervals can vary depending upon the specific network and the information that needs to be analyzed.

[0057] Referring now to block 212, the poller establishes a connection to each gateway 110 in the network 100. For example, this connection can be established using SNMP (Simple Network Management Protocol) or CLI (Command Line Interface). After a connection is established, operating parameter data is requested by the poller and transmitted from each gateway, as indicated by block 214. In one embodiment, a connection is established with each gateway and information is gathered from that gateway before the next gateway is contacted. In another embodiment, connections are open with multiple gateways at the same time.

[0058] Once the data has been received, it is written to a raw data file as indicated by block 216. In the preferred embodiment, two files are generated. A first file includes the raw data as received from the gateways. The second file is a summary file that can be updated after each polling interval in the reporting period. After each reporting period, a new summary file can be created. In addition, multiple summary files can be created for different but overlapping reporting periods.

[0059] At the end of each reporting period, the data is analyzed to create a health report.

This step is shown in block 218. For example, the reporting period can be daily so that a script is set up to automatically operate at a particular time (e.g., midnight GMT). The health report includes a summary of the information gathered from the operating parameter data. Further details on the health reports are provided below.

[0060] As indicated by step 219, the health report can be generated as a file and transmitted to recipients. In the preferred embodiment, the health report is formatted in ASCII text and automatically mailed using a UNIX™ mail utility, which sends the report to a preconfigured list of recipients. The recipients are typically network maintenance and support staff who have the responsibility of monitoring the stability of the network so that failures can be avoided.

[0061] Figure 4 provides a flowchart 300 that illustrates the process from the perspective of the support staff. The health report is received, as indicated by block 310. This report is preferably received via e-mail on a daily basis. Due to the global nature of many networks, different staff will receive the same report at different times depending upon their location (e.g., time zone).

[0062] In an alternate embodiment, the system can generate multiple health reports that cover overlapping time periods. For example, staff in Tokyo can receive a report at 6 am local time. Nine hours later, an updated report can be generated and sent to staff in London, where it is 6 am local time. Eight hours later, yet another updated report can be generated and sent to staff in Los Angeles, where it is now 6 am local time. Each updated reports can, but does not need to be, sent to staff worldwide. The number of updates within a reporting period can vary but a typical case would include between two and four updates (e.g., every 6 to 12 hours).

[0063] As indicated by block 312, the staff will review the report to check for flags. As will be discussed below, the reporting system can be programmed to compare the operating parameters to thresholds determined from a stable network. A flag (or flags) will be automatically generated when any of the parameters is outside this predetermined threshold.

[0064] Another feature supported by the system is trend analysis. Certain items may appear to be operating appropriately in any given chart but negative trends can be determined by monitoring changes over time. Accordingly, some of the parameters can be saved daily, monthly or yearly to view trends.

[0065] Based on the report, the staff can determine if attention is required, as indicated by block 314. If no attention is required, the report can be archived and the process is complete for the reporting period (block 318). If attention is required, the data stored in the raw data and summary files can be consulted. Based on this data, the cause of the potential instability can be determined and appropriate corrective action be taken. In this manner, instabilities are predicted and corrected before adversely affecting the network.

[0066] Examples of the specific parameters that can be monitored will now be described. Tables 1 provides example of a health reports that can be generated to monitor the stability of the gateway 110. This health report focuses on a number of parameters. Each parameter will be discussed in turn.

Table 1

SECURE INTERWORKING GATEWAY (SIG) NETWORK DAILY HEALTH REPORT

Report For: March xx, 2004

All times reported are in GMT.

=====

DEAD IKE SAs

=====

LOCATION	NODE	TOTAL IKE SAs	DEAD IKE SAs	% DEAD IKE SAs
Tokyo	tkyxbsnz	x	x	x.x
Manhattan	mnhtsbsnd	x	x	x.x
Richardson	rcdsdgtwp	x	x	x.x

=====

NODE THROUGHPUT

=====

Total Remote Access Customers: xx

Total Nodes: x

Total Bytes Received: xxxxxxx Kbytes/day

Total Bytes Sent: xxxxxxx Kbytes/day

LOCATION	NODE	RECEIVED	SENT
		KBYTES	KBYTES
Tokyo	tkyxbsnz	xxxx	xxxx
Manhattan	mnhtsbsnd	xxxxxx	xxxxxx

Richardson rcdsdgtwp xxxxx xxxxx

NODE CONFIGURATION

Total number of NODES: x

Limits: MAX VPRNs: xxxx

MAX CONNECTIONS PER SSC: xxxx

MAX IPSEC PER SSC: xxxx

Chassis Thresholds: VPRN WARN: xx%

VPRN AUGMENT: xx%

CONNECTIONS WARN: xx%

CONNECTIONS AUGMENT: xx%

IPSEC WARN: xx%

IPSEC AUGMENT: xx%

Total number of Nodes with any value above its threshold: x

FLAG NOTES: ! = value over a threshold

WARN = Warning - getting close to limit

AUG = Node needs augmenting

Tokyo	tkyxbsnz	xx	xx	x	x	x.x
Manhattan	mnhtsbsnd	xx	xxx	x	xx	x.x
Richardson	rcdsdgtwp	x	x	x	x	x.x

=====

CMC (Control and Management Card) TOGGLERS

Total Number of CMC toggles: x

=====

SFC (Switch Fabric Card) TOGGLERS

Total Number of SFC toggles: x

=====

CMC CPU UTILIZATION

Average of all Peak CMC CPU Utilizations: x.xx%

Peak CMC Px Utilization: x.xx% in Manhattan on mnhtsbsnd.xx at xx:xx

Peak CMC Px Utilization: x.xx% in Tokyo on tkyxbsnz.xx at xx:xx

CMC CPU Utilization Threshold: xx %

Number of nodes with either utilization over threshold: x

LOCATION	SWITCH	PEAK %	AVG %	PEAK %	AVG %	FLAG
		CMC Px	CMC Px	CMC Px	CMC Px	
		UTIL	UTIL	UTIL	UTIL	
-----	-----	-----	-----	-----	-----	-----
Tokyo	tkyxbsnz	x.xx	x.xx	x.xx	x.xx	
Manhattan	mnhtsbsnd	x.xx	x.xx	x.xx	x.xx	
Richardson	rcdsdgtwp	x.xx	x.xx	x.xx	x.xx	

=====

CMC MEMORY UTILIZATION

=====

Average of all Peak CMC Memory Utilizations: xx.xx%

Peak CMC Memory Utilization: xx.xx% in Manhattan on mnhtsbsnd at xx:xx

CMC Memory Utilization Threshold: xx%

Number of nodes with utilization over threshold: x

LOCATION	SWITCH	PEAK %	AVG %	FLAG
		CMC MEM	CMC MEM	
		UTIL	UTIL	
-----	-----	-----	-----	-----
Tokyo	tkyxbsnz	xx.xx	xx.xx	
Manhattan	mnhtsbsnd	xx.xx	xx.xx	
Richardson	rcdsdgtwp	xx.xx	xx.xx	

SSC (Subscriber Services Card) RESETS

=====

Total Number of card resets: x

=====

NO STANDBY CMC

=====

Total switches without a standby CMC: x

=====

NO STANDBY SFC

=====

Total switches without a standby SFC: x

=====

SSPx (Subscriber Services Processor) UTILIZATION

=====

Average of all Peak SSPx CPU Utilizations: xx.xx%

Peak SSPx CPU Util: xx.xx% in Tokyo on tkyxbsnz.x.x.x at xx:xx

Average of all Peak SSPx Memory Utilizations: x.xx%

Peak SSPx Memory Util: x.xx% in on at :

Limits: SSPx FIBs: xxxx

SSPx CONNECTIONs: xxx

Thresholds: SSPx CPU Utilization Threshold: xx%

SSPx Memory Threshold: xx%

SSPx FIB Threshold: xx%

SSPx Connection Threshold: xx%

LOCATION	NODE.SSC.SSM.SSP	MAX % UTIL		MAX	
		CPU	MEMORY	FIBs	CONNs
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Tokyo	tkyxbsnz.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Manhattan	mnhtsbsnd.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x

Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x
Richardson	rcdsdgtwp.x.x.x	xx	x	x	x

[0067] The first entry in the health report relates to dead IKE SAs. IKE SA stands for Internet Key Exchange Security Association, an object that is created when an IPSec tunnel is established. The gateway 110 has been known to fail to clean up IKE SAs when IPSec tunnels are disconnected, resulting in an accumulation of these objects. IKE SAs can be counted, and active tunnels can be counted, and the difference is reported as Dead IKE SAs. An incrementing count serves notice that the disconnect process is not functioning properly, so that maintenance activity can be scheduled.

[0068] The next entries relate to node throughput and node configuration. The node throughput measures the amount of traffic through each node. The VPRN identifier contains a unique name for each customer and service. From a list of all VPRNs these unique names are counted to provide the number of customers for a particular service. This list includes internal and external, billable and non-billable customers. The number of nodes is derived from a manually maintained list. Received and Sent KBytes are totaled for all trunk and access interfaces, and include externally received packets, internally generated packets, signaling packets and test packets.

[0069] Node level limits exist for NUMBER OF VPRNs, NUMBER OF CONNs and MAX IPSEC. In this context, VPRN stands for Virtual Private Routed Network, the term used by

Nortel for a VPN (Virtual Private Network), as well as for a Virtual Router. In this section, VPRN refers to a Virtual Router in a particular gateway. CONNs stand for Connections, the generic attachments to the gateway. Connections comprise PVCs, trunks, and IPsec Tunnels. There is a higher limit for Connections than for IPsec tunnels.

[0070] For VPRNs the node level limit is stated directly. For NUMBER OF CONNs and MAX IPSEC, the stated per SSC (subscriber service card) limit is multiplied by the number of SSCs to determine the Node limit. Two thresholds are stated as a percentage of the Node limit: WARN and AUGMENT. Any metric exceeding the WARN threshold will be flagged (!) after the reported value, and the FLAG field will read WARN, indicating that the metric or metrics should be monitored carefully. Any metric that also exceeds the AUG threshold changes the FLAG field to read AUG, indicating that capacity should be augmented.

[0071] The next entries in the report relate to CMC (control and management card) and SFC (switch fabric card) toggles. In each gateway there is an active and a standby CMC. If the active CMC becomes inactive for any reason causing the standby CMC to assume the active role, a Toggle is said to have occurred. The Number of CMC Toggles is a count of BSNs that have experienced one or more CMC toggles in the reporting period.

[0072] In each gateway there is also an active and a standby SFC. If the active SFC becomes inactive for any reason causing the standby SFC to assume the active role, a Toggle is said to have occurred. The Number of SFC Toggles is a count of BSNs that have experienced one or more SFC toggles in the reporting period.

[0073] The CMC CPU utilization section provides a summary of the utilization percentages of the two processors on the CMC, responsible for system wide routing protocols, FIB (forwarding information base) generation, and system configuration and management functions.

The highest utilization is reported as Peak, along with the average of all the samples. A flag is set if any utilization exceeds the listed threshold.

[0074] CMC memory utilization is also monitored and reported. Memory is a critical system resource on the Control and Management Card (CMC), which performs all the routing functions in the system, including routing protocol messaging, creation of Routing Information Bases (RIBs) for all the Virtual Routers, and the creation and distribution of Forwarding Information Bases (FIBs) to the SSCs (Subscriber Services Cards). The CMCs do not forward traffic, so this memory does not comprise buffers, but is the resource for total route capacity in the System. Utilizations over the listed threshold are flagged.

[0075] The number of gateways that have experienced one or more resets on one or more of the SSCs in the chassis during the reporting period is tracked in the section entitled SSC card reset. Likewise, the number of gateways that do not have a functional standby CMC is reported in the No Standby CMC section and the number of BSNs that do not have a functional standby SFC is reported in the No Standby SFC section.

[0076] The next section in the health report relates to SSP4 (Subscriber Services Processor) utilization. The Subscriber Services Card (SSC) does the forwarding for the gateway. Each SSC has four SSMs (Subscriber Services Modules), and each SSM in turn has four SSPs (Subscriber Services Processors). There are four SSP4s per SSC. One of these SSPs, the SSP4, is the CPU that processes IPsec tunnels. Since most Remote Access connections are IPsec, SSP4 Utilization is a direct measure of gateway system forwarding capacity. For each SSP4 in the gateway, this section reports Max CPU Utilization %, Max Memory Utilization %, Max FIB count, and Max Connection count. A FIB is a Forwarding Information Base, the forwarding table, one per Virtual Router. The individual SSP4s are identified by <switchname>.<slot>.

<module>.4. For all values, the peak readings during the reporting period are shown. A flag is set if any of the values exceeds the stated threshold.

[0077] Another parameter that should be monitored relates to flowcache. The flowcache is a specialized piece of memory that is used to store current connection details. The most elementary of these is a flow, which is a unique 5-tuple of Source IP Address, Source Port Number, Destination IP Address, Destination Port Number, and Protocol. When the cache contains two such flows for the same Protocol, such that the Source IP Address and Source Port Number of the first flow are the same as the Destination IP Address and Destination Port Number of the second flow, and the Destination IP Address and Destination Port Number of the first flow are the same as the Source IP Address and Source Port Number of the second flow, these two flows are logically grouped together within the flowcache and categorized as a 'connection'. Similarly, when multiple connections exist in support of a single higher layer protocol, (HTTP, SIP, Real Audio, etc.), they are logically grouped together within the flowcache to form a 'conversation'. The flowcache has a number of uses. Gateway services and policies applied to packets associated with connections and conversations typically vary. The first time a packet associated with a particular connection or conversation is forwarded, the service and policy list is looked up and stored in the flowcache. From that point onward, successive packets associated with the same connection or conversation are processed with the associated service and policy list in the flowcache, as long as the connection or conversation is active. This eliminates the need to look up the service and policy list for every packet that is forwarded, resulting in significant CPU cycle savings.

[0078] Another use of the flowcache is for state-aware firewall packet filtering. For example, flows originating from the secure side of the firewall outbound can be allowed, but

flows originating from the non-secure side of the firewall inbound can be disallowed, unless such an inbound flow can be associated with an outbound flow to form a connection, in which case it is allowed, being a response to a connection initiated from the secure side of the firewall. Once this connection is cached, packets associated with it can flow both inbound and outbound, as long as the connection is active. Once the connection is inactivated, inbound packets with the same 5-tuple are no longer allowed to pass through the 'pin hole'.

[0079] While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications and combinations of the illustrative embodiments, as well as other embodiments of the invention, will be apparent to persons skilled in the art upon reference to the description. It is therefore intended that the appended claims encompass any such modifications or embodiments.